



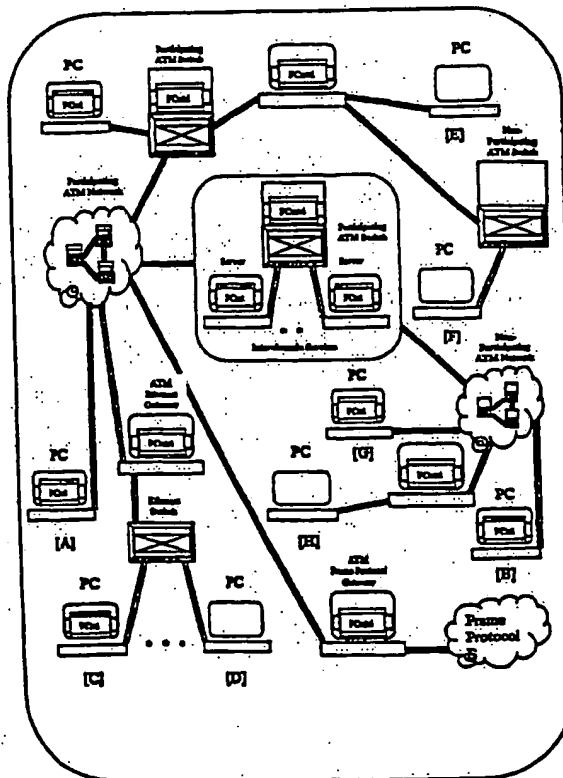
INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁷ : H04L 12/56, 12/28	A1	(11) International Publication Number: WO 00/62496 (43) International Publication Date: 19 October 2000 (19.10.00)
(21) International Application Number: PCT/US00/09964 (22) International Filing Date: 13 April 2000 (13.04.00) (30) Priority Data: 09/291,485 14 April 1999 (14.04.99) US (71)(72) Applicant and Inventor: TREBES, Harold, Herman, Jr. [US/US]; 1361 Middlesex Avenue N.E., Atlanta, GA 30306 (US). (74) Agent: GOLDMAN, Joel, S.; Troutman Sanders LLP, Suite 5200, 600 Peachtree Street N.E., Atlanta, GA 30308-2216 (US).	(81) Designated States: AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CR, CU, CZ, DE, DK, DM, DZ, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, UZ, VN, YU, ZA, ZW, ARIPO patent (GH, GM, KE, LS, MW, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG). Published With international search report.	

(54) Title: A SYSTEM AND METHOD FOR PROVIDING PEER-ORIENTED CONTROL OF TELECOMMUNICATIONS SERVICES

(57) Abstract

In a telecommunications network environment including non-participating elements and participating elements, a method for providing a telecommunications service between a first peer element connected to the telecommunications network environment and a second peer element connected to the telecommunications network. At a first peer element, an indication of the type of telecommunications service to be provided between the first peer element and the second peer element is received. A telecommunications service template in association with the indicated telecommunications service is determined, the telecommunications service template including instructions for configuring the non-participating elements of the telecommunications network environment to provide the indicated telecommunications service and instructions for configuring the participating elements of the telecommunications network environment. The telecommunications service template may further comprise routing instructions for the non-participating elements of the telecommunications network environment and routing instructions for the participating elements of the telecommunications network environment.



FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece			TR	Turkey
BG	Bulgaria	HU	Hungary	ML	Mali	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MN	Mongolia	UA	Ukraine
BR	Brazil	IL	Israel	MR	Mauritania	UG	Uganda
BY	Belarus	IS	Iceland	MW	Malawi	US	United States of America
CA	Canada	IT	Italy	MX	Mexico	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NE	Niger	VN	Viet Nam
CG	Congo	KE	Kenya	NL	Netherlands	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NO	Norway	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	NZ	New Zealand		
CM	Cameroon	KR	Republic of Korea	PL	Poland		
CN	China	KZ	Kazakhstan	PT	Portugal		
CU	Cuba	LC	Saint Lucia	RO	Romania		
CZ	Czech Republic	LI	Liechtenstein	RU	Russian Federation		
DE	Germany	LK	Sri Lanka	SD	Sudan		
DK	Denmark	LR	Liberia	SE	Sweden		
EE	Estonia			SG	Singapore		

2 **A SYSTEM AND METHOD FOR PROVIDING PEER-
ORIENTED CONTROL OF TELECOMMUNICATIONS
SERVICES**

4 **REFERENCE TO PROVISIONAL APPLICATION**

 This non-provisional patent application claims the benefit of provisional patent
6 application No. 60/081,710 filed on April 14, 1998 and entitled "Peer-Oriented Control and
Service Creation in a Internetworking Environment", which is incorporated by reference
8 herein.

FIELD OF THE INVENTION

10 This invention relates to telecommunications and, more particularly, to a
system and method for providing peer-oriented control of telecommunications services
12 through the use of an application level or "logical level" control mechanism.

BACKGROUND

14

 The deregulation of telephone companies, or Telcos, has lead to increased
16 competition. In many cases, Telcos and other carriers, are freed by statute to offer to any

other competitor with carrier status, substantial discounts designed to level the competitive
2 playing field for the network access or bandwidth delivery portion of the access market.
Then more bandwidth is available at a cheaper price than was previously possible. However,
4 bandwidth and services are often bundled together and sold. Thus, many of the savings and
benefits of the cheaper bandwidth are not realized by the user.

6

Thus, there is a need for an application level or logical level control
8 mechanism for communication services used in support of various peer-oriented types of
applications. There is also a need for a control mechanism that is orthogonal to the
10 underlying native control mechanisms of the network being used. In other words, the control
mechanism would function regardless of the intervening control mechanisms of the network.
12 This capability allows application developers to use network services as components of their
applications with minimal concern for the implementation of those services. Thus, the
14 cheaper bandwidth may be purchased from telcos, without the added costs of attached
services.

16

These needs will become apparent from the following text. For a number of
18 years now, telecommunications and networking have been assuming increasingly strategic
roles supporting the fundamental structure and operation of companies. One milestone that
20 may be noted on this evolutionary path comes from an article in Business Week magazine
published in the issue of February 8, 1993. This article entitled "The Virtual Corporation"
22 popularized the discussion of management concepts and practices that had been discussed in
organizational management literature and practiced to varying degrees by organizationally
24 sophisticated companies for some time. The introductory comments on the topic printed on
the magazine's cover provided the framework for considering the topic:

2 Big, complex companies usually can't react fast enough. Small, nimble ones
may not have the muscle. What's the answer? A new model that uses
4 technology to link people, assets, and ideas in a temporary organization. After
the business is done, it disbands. It's called the virtual corporation. Just
6 another management fad -- or a vision of the future?

8 Although the seeds of recent telecom and networking phenomenon are present
within these introductory words, the current explosion of technologies, products, and
10 variations for their strategic and tactical use was not fully foreseen or understood or at least
was not expressed at this point in time.

12
In today's business environment, there is not so much of a revolution as there
14 is a super accelerated evolution in the economic and information fabric in which business
operates. Information accessibility and electronic connectivity combine to provide the
16 equalizer on the frontier of global business and economic opportunity. As communication
and networking technology developers seek to keep up with the escalating demand for more
18 dynamic and easier communication capabilities, there is a shift in their market orientation.
This shift in technology providers' approach to their market may be viewed as an indication
20 of underlying environmental forces which will favor significant architectural changes in the
structure of networks and the mode of creation for network services supporting collaborative
22 applications.

24 In looking at the primary market approach, two fundamental orientations can
be identified: the technology push and the application pull. The technology push orientation

says "tell us what your network-related problems are and we will show how to solve them
2 using a set of products." The application pull says "here are application level solutions to
problems that your business currently has or is likely to have based on the evolution of the
4 business environment and this solution is currently implemented using a set of products." The
technology push is traditionally associated with manufacturers and generic networking
6 resellers and integrators, whereas the application pull is normally associated with the true
vertical market specialists.

8

Projecting the technology push and application pull orientation into the
10 solution mindset of target market potential customers highlights the two corresponding
dominant customer orientations: network centric (associated with the technology push
12 orientation) and application centric (associated with the application pull orientation). Telcos
and Competitive Access Providers (CAPS) can be used to illustrate these points for network
14 resource providers, but it is important to realize that similar distinctions exist within end-user
organizations where the network support organization generally holds network centric views,
16 whereas the operating business units generally hold application centric views.

18

The network supplier market as represented by Telcos and CAPS, especially in
the U.S., provides an interesting example with which to illustrate significant aspects of these
20 differing orientations to the potential customer solution evaluation process. Since the start of
deregulation and the opening of network access to competitive pressures, there has been an
22 evolutionary force, i.e. competition, at work on the structure and basic business positioning of
Telcos and businesses that would compete against them. Prior to the start of open
24 competition for the network access market, the Telcos, as well as their limited competitors the
CAPS, can be characterized as holding primarily what has been called network centric views

of solutions. Characteristic of this view is the bundling of service and feature differentiators with combinations of "raw" bandwidth delivery infrastructures to create "products" which would be sold in a manner consistent with the technology push orientation. A case can be made that much the same situation currently exists within the network support groups that currently support the network infrastructures upon which the applications of large, medium and increasingly small companies are deployed.

With the coming of open competition in the network access markets, however, pressures of the new business environment have caused a fundamental shift in the structure and the business approach of such organizations. Specifically, what had been the service and non-connectivity related features of the "product" (aggregately identified as the "product differentiators") are rapidly being separated from the bandwidth delivery infrastructure and moved into non-regulated business units that function at the retail level and which compete with resellers, network integrators, and vertical market specialists. This evolution has been caused largely by new regulatory statutes that force the Telcos, or any other carrier, to offer to any other competitor with carrier status, substantial discounts designed to level the competitive playing field for the network access or bandwidth delivery portion of the access market.

This business environment situation has started an irreversible shift in the value creation chain for telecommunication services in which the biggest "added value" link will shift from the "wires" business associated with bandwidth transmission and delivery to the "product differentiator" services and features. The monolithic product set once associated with the telecommunications industry has been split into an interoperable bandwidth transport and delivery access infrastructure commodity and a separate service/feature creation

opportunity that has significant potential for differentiation and value creation. This
2 evolutionary transformation, which is now underway, has significant implication for the
marketing channel mix of networking product vendors as the relative importance of
4 technology push versus application pull orientations seek a new equilibrium in the new
business environment.

6

One of the bandwidth transmission mediums is the asynchronous transfer
8 mode, or ATM, transmission medium. The current service creation and network control
architectures fail to adequately harness the potential flexibility of the ATM transport
10 mechanism. The potential to carry any type of traffic, along with the ability to link
terminating points over a mixture of public and private network resources in an on-demand
12 fashion, opens up a whole new realm of technological challenges and economic potential, the
ramifications of which are only beginning to be grasped.

14

However, despite the available bandwidth, there is still a need for the ability of
16 individual end-users, or peers, to have the ability to set up and control services that have been
typically set up and controlled by the telcos.

18

There are at least three potential reasons which might help to explain why this
20 need has not already been met. First, a reliable, distributed, peer-oriented service creation
facility is more difficult to develop as compared with currently existing telecommunications
22 service creation mechanisms. Existing mechanisms may be viewed as utilizing a client-
server model in the sense that a session requests a certain service capability from the network
24 and the network control function responds by determining if the resources are available and
then signaling to switches to establish the service. Part of the added complexity for a peer-

oriented mechanism comes from the use of active peer-agents negotiating to set up and
2 maintain a requested service. Some of the factors involved which contribute to this additional
complexity include problems of managing distributed threads of control, including problems
4 of process synchronization, as well as the greater risk of message loss or corruption
introduced through the increased use of communication links connecting the collaborating
6 processes which greatly increases the need for additional fault detection mechanisms.

8 The second reason which may contribute to the lack of such a solution
concerns the evolution and current state of the public telecommunications networks.
10 Metropolitan and wide-area networks are generally established utilizing the physical facilities
of public telecommunications carriers. The networks that these carriers have deployed have
12 evolved from networks which were originally established to handle analog voice traffic
through switched circuit technology. A case can be made that much of the current
14 architecture for service creation has come to be as the result of incremental response to
evolutionary trends in service needs and resource capabilities as well as the cost structures
16 that were associated with possible development path options. The development and
introduction of ATM has provided the first standards-based transport mechanism that is
18 designed to support all types of traffic. When combined with the User Network Interface
(UNI) Staff ATM Forum (1995) and the Private Network to Network Interface Staff ATM
20 Forum (1995) developed by the ATM Forum, a case can be made that the basis for an
alternative user-controlled network service creation and control paradigm has been created.
22 So, the second reason such work has not been performed is that there was no compelling
reason to undertake what is a much more difficult architecture to design and implement as
24 long as the network was predominantly a circuit-switched infrastructure.

2 The third reason concerns the growth of capabilities in the areas of both
hardware and software. In order to develop a service creation process which utilizes a
4 distributed architecture functioning in a real-time collaborative mode, great demands are
placed on the hardware and software system components. An observation might be made that
6 the rapidly dropping cost of processing power, along with the advances in methodologies,
CASE tools and the development of middleware platforms are enabling factors that needed to
8 be available before distributed systems approaches to communications infrastructure could go
forward on a commercial scale. Therefore, the third factor which might be considered as
10 hindering similar research in the past is the potentially diminished interest due to the
inadequacy of the commercial tools and techniques then available.

12

 Thus, there is a need for an application level or logical level control
14 mechanism for communication services used in support of various peer-oriented types of
applications. There is also a need for a control mechanism that is orthogonal to the
16 underlying native control mechanisms of the network being used. In other words, the control
mechanism would function regardless of the intervening control mechanisms of the network.
18 This capability allows application developers to use network services as components of their
applications with minimal concern for the implementation of those services.

20

SUMMARY

22

 The present invention meets the above-described needs by extending the core
24 networking technology more directly into the world of the application, thereby providing a

network-aligned infrastructure that is capable of better supporting the development and
2 deployment of collaborative applications. Embodiments of the present invention allow an
end-user to control creation of telecommunications services from the edge of the
4 telecommunications network. Previously, telecommunications services have been created
within the network, such as by the carriers or telcos.

6

In one aspect, the present invention is a method, in a telecommunications
8 network environment including non-participating elements and participating elements, for
providing a telecommunications service between a first peer element connected to the
10 telecommunications network environment and a second peer element connected to the
telecommunications network. At a first peer element, an indication of the type of
12 telecommunications service to be provided between the first peer element and the second peer
element is received. A telecommunications service template in association with the indicated
14 telecommunications service is determined, the telecommunications service template including
instructions for configuring the non-participating elements of the telecommunications
16 network environment to provide the indicated telecommunications service and instructions
for configuring the participating elements of the telecommunications network environment.
18 The telecommunications service template may further comprise routing instructions for the
non-participating elements of the telecommunications network environment and routing
20 instructions for the participating elements of the telecommunications network environment.
The instructions to configure the participating elements and non-participating elements of the
22 telecommunications network environment are executed to provide the telecommunications
service. Data between the first peer element and the second peer element is transmitted via a
24 predefined transmission protocol indicated by the telecommunications service template, the
data including the routing instructions for the non-participating elements of the

telecommunications network environment in a header portion of the predefined transmission
2 protocol and the routing instructions for the participating elements of the telecommunications
network environment in a payload portion of the predefined transmission protocol.

4

In one aspect, the present invention allows a user to set up a
6 telecommunications service at the edge of a network. Thus, bandwidth may be purchased at a
discount and the bandwidth may be allocated to the services defined by the user. These
8 services are created by the user rather than being created by the carrier and sold in a bundle
with the bandwidth. The present invention may function with both participating and non-
10 participating networks. The participating networks include active elements that route the data
based upon instructions including in the payload and/or control portion of ATM. Non-
12 participating networks route the ATM cells without disturbing the present invention. Thus,
the present invention is not limited by participating networks. The present invention is also
14 useful as an encoding or encrypting mechanism because the data is transmitted from one peer
element and then decoded by a second peer element. Thus, encoding and encrypting is an
16 inexpensive and useful feature of the present invention. The present invention also includes
active participating network elements that include such useful features as self-healing and
18 communication with each other.

20 BRIEF DESCRIPTION OF THE DRAWINGS

22

Fig. 1 is a high-level illustration of an embodiment of the present invention.

2 Figs. 2A-2C are N-square charts for the workstations, servers, and
participating ATM switches of an embodiment of the present invention.

4 Fig. 3 is an illustration of the participating and non-participating boundaries of
the infrastructure of an embodiment of the present invention.

6 Fig. 4 is an illustration of high-level use cases for the infrastructure of an
8 embodiment of the present invention.

10 Fig. 5 is a description of the use case actors.

12 Figs. 6A-6B are summaries of the participation of the actors in the various use
cases.

14 Figs. 7A, 7B, 8A, 8B, 8C, 9A, 9B, 9C, and 9D are illustrations of use cases
16 involving an embodiment of the present invention.

18 Fig. 10 is an illustration of system function classification categories.

20 Figs. 11A-11C, 12A-12C and 13A-13C are illustrations of the functions
associated with the workstations, participating switches and domain services server of an
22 embodiment of the present invention.

24 Fig. 14 is an illustration of the Zachman framework for the peer-oriented
infrastructure of an embodiment of the present invention.

2 Figs. 15 and 16 are illustrations describing architectural patterns.

4 Fig. 17 is an illustration of the control pattern of an embodiment of the present
invention.

6 Figs. 18A and 18B are lists of realization mechanisms.

8 Fig. 19 is a Quality Function Deployment Process diagram.

10 Fig. 20 is a QFD spreadsheet for peer-controlled infrastructure.

12 Fig. 21 is a QDS spreadsheet for peer-controlled infrastructure.

14 Figs. 22A and 22B are illustrations of the raw application scoring for the
16 alternatives listed at the bottom of Fig. 21/

18 Fig. 23 is an illustration summarizing the applicability of major functional
elements to the major architectural components of the infrastructure.

20 Fig. 24 is an illustration of the impact of realization mechanisms on the
22 functional classes.

24 Figs. 25 and 26 are illustrations of the buildup of cumulative technical impact
associated with the successive selection of available realization mechanisms.

2 Figs. 27 and 28 are illustrations of the successive redistribution of technical
impact among the function system areas as each successive choice of realization mechanisms
4 is added to the architectural framework.

6 Fig. 29 is a deployment model for the peer-oriented infrastructure of an
embodiment of the present invention.

8 Fig. 30 is an illustration of the peer-control approach to reconciling
10 networking perspectives.

12 Figs. 31A-31B are flowcharts illustrating the user interaction with a user
interface to set up a telecommunications service in accordance with an embodiment of the
14 present invention.

16 Fig. 32 is a flowchart for establishing a call between peer elements in
accordance with an embodiment of the present invention.

18 Fig. 33 is a diagram of the metanetwork capabilities of an embodiment of the
20 present invention.

DETAILED DESCRIPTION

2

In one aspect, the present invention is a method for providing peer-oriented control of a telecommunications and data networking-based collaborative service. The present invention is concerned with the structure of an application level or "logical level" control mechanism for communication services used in support of various peer-oriented types of applications. The present invention is orthogonal to the underlying native control mechanisms of the network being used. This capability allows application developers to use network services as components of their applications with minimal concern for the implementation of those services.

12

The target platform for this infrastructure is based on cell and packet based networks utilizing a virtual circuit type implementation mechanism, although implementation over a datagram type implementation mechanism or other mechanisms is possible. The present invention focuses on the Asynchronous Transfer Mode (ATM) communications platform although it may be extended to other cell and packet based communication infrastructures.

18

The physical characteristics of the ATM communications environment, along with the nature of a peer-oriented service creation process, lead to the need for a solution domain based on distributed, cooperative processing. The peer-oriented service creation process is viewed as a collaborative goal-seeking activity among network resource elements which may continue only up to the point that the service is created (this would parallel the current service creation model used in telecommunications networks) or preferably would

24

continue functioning in a collaborative goal-seeking fashion throughout the duration of the
2 session utilizing the service.

4 There are a whole range of issues associated with a distributed processing
solution. The following issue areas are indicative but not exhaustive of issues that are
6 significant in considering a peer-oriented service creation paradigm:

- 8 • Scalability of solution approach
- 10 • Impact of competing architecture and mechanism design approaches
- 12 • The need for infrastructure to precede applications
- 14 • Impact of end-user versus resource owner issues in cross-domain
16 problems
- 18 • Economic feasibility of design approaches given the large installed
base of equipment.

20 Before addressing the foregoing issues and proceeding with a more detailed
description of the present invention, some important terms used herein are defined below.

22

Definitions of Terms

2

Orthogonal Control

4

This is a concept being developed by the present invention. The basic notion behind orthogonal control mechanisms is to make the network resources transparent to the collaborative application environment. Going beyond the notions of a simplifying API generally provided by middleware, the present invention defines orthogonal control as infrastructure architectural mechanisms that transparently translate or map control notions related to collaborative efforts directly into supporting control mechanisms for the network.

12

Quality of Service (QoS)

14

Quality of service is a term which refers to the set of ATM performance parameters that characterize the traffic over a given virtual connection (VC). These parameters include the CLR (cell loss ratio), CER (cell error rate), CMR (cell misinsertion rate), CDV (cell delay variation), CTD (cell transfer delay), and the average cell transfer delay. Five service classes have been defined by the ATM Forum in terms of QoS parameters. There is a correlation between these classes and the ATM Adaption Layers defined later. The QoS service classes are:

22

Class 0 - best efforts service

24

Class 1 - specifies the parameters for circuit emulation - associated with AAL1

2 Class 2 - specifies the parameters for VBR audio & video - associated with

AAL2

4

6 Class 3 - specifies the parameters for connection-oriented services - associated
with AAL3/4 and AAL5

8

Class 4 - specifies the parameters for connectionless data transfer - associated
with AAL3/4 and AAL5

10

ATM Adaption Layer (AAL)

12

14 The ATM adaption layer is a collection of standardized protocols that provide
services to higher layers by adapting user traffic to a cell format. The AAL is divided into the
convergence sublayer (CS) and the segmentation and reassemble (SAR) sublayer. The four
16 AAL types currently defined are:

18

AAL1 - a protocol standard used for the transport of constant bit rate (CBR)
traffic (i.e., audio and video) and for emulating TDM-based circuits (i.e., DS 1, E1, etc.).

20

AAL2 - a protocol standard for supporting time-dependent variable bit rate
22 (VBRRT) connection-oriented traffic (i.e., packetized video and audio).

2 AAL3/4 - AAL type 3 and 4 provide a protocol standard for supporting both
connectionless and connection-oriented variable bitrate (VBR) traffic. This AAL is also used
to support SMDS (switched multimegabit digital service).

4

AAL5 - a protocol standard for supporting the transport of lightweight variable
6 bit rate (VBR) traffic and signaling messages. This AAL is also used to support frame relay
services.

8

Service Creation

10

A service in the telephony and data networking world is generally taken to
12 mean a defined action which creates a facility (i.e., a telephone call) or performs a function
(i.e., forwards a call) performed by network control elements in response to a request by a
14 subscriber or user. Service creation may be defined as the complete process of IN (intelligent
network) service creation, including design, specification, development, and verification.
16 Within this application, however, a slightly different notion will be used. That is, service
creation will be used to mean the act or functioning of control elements and network
18 resources together in order to establish the facility or perform the function requested by the
subscriber or user.

2 **Peer**

4 As used herein, peers are any two or more end units that want to collaborate
together in a service. Peers, unless they are on a LAN, must communicate with one another
6 through some public facility. A common way is TCP/IP Internet protocols.

8 **Peer-oriented Service Creation**

10 This application builds on the previous definition for service creation with the
following definition for peer-oriented service creation. The term "peer-oriented service
12 creation" is defined as a control paradigm that is based on separating some amount of control
or use interpretation from the network and assigning that control to participating workstations
14 which are being used to support the user interface for a collaborative infrastructure. Applying
this notion, participating workstations are used to request bandwidth with a specified quality
16 of service to connect participants but the workstation, in collaboration with one another, also
supply a collaborative control environment and the additional information resources
18 necessary to establish the facility or perform the function (i.e., provide the service) that is
desired.

20

Active Networks

22

 Active networks are a novel approach to network architecture in which the
24 switches of the network perform customized computations on the messages flowing through
them. This approach is motivated by both lead user applications, which perform user-driven

computation at nodes within the network today, and the emergence of mobile code
2 technologies that make dynamic network service innovation attainable. With regard to the
present invention, the notion just expressed is explicitly enlarged to include similar behavior
4 at the network interface level (network interface card - NIC) of workstations and servers
attached to, and functioning as part of, the larger notion of network which is an infrastructure
6 connecting collaborators. In an active network, elements, such as switches, obtain
information about the status of the network and circulate this information throughout the
8 network for self-healing, controlling traffic flow, etc.

10 Edge Networks

12 The approach adopted in this application focuses on the use of ATM as the
foundation of the network used to support the collaborative environment. As various LAN
14 technologies such as Ethernet and Token Ring are currently the dominant network platform
for supporting collaborative efforts, there must be a method of interfacing these technologies
16 to the ATM infrastructure. Viewing ATM networks as the core of a new infrastructure places
other networking technologies at the "edge" of the network. It is at the boundary between the
18 ATM core network and the other networking technologies at the edge that issues of protocol
and control translation become significant.

20

ATM Forum Network Reference Model

22

The Network Reference Model of the ATM Forum extends the model
24 developed by the ITU-T by taking care to distinguish between the private and public parts of
an ATM network. The model serves to identify the following key interfaces described below:

2 User-Network Interface (UNI) -- User network interface. The interface
defined as a set of protocols and traffic characteristics (i.e., cell structure) between the CPE
4 (user) and the ATM network (ATM switch). The ATM Forum specifications refer to two
standards being developed, one between a user and a public ATM network, called public
6 UNI, and one between a user and a private ATM network, called P-UNI.

8 Private [Network-Node or Network-Network] Interface (PNNI) -- PNNI is a
switch-to-switch protocol developed within the ATM Forum to support efficient, dynamic,
10 and scalable routing of SVC (switched virtual circuit) requests in a multivendor private ATM
environment.

12 Broadband Inter-Carrier Interface (B-ICI) -- The broadband intercarrier
14 interface is a specification that enables two adjacent public ATM networks to interconnect
and provide a set of end-to-end services.

16

Cells In Frames

18

Cells In Frames (CIF) is ATM with variable length packets on the lines and
20 trunks. The CIF Alliance has specified a protocol which allows ATM to be embedded into
various frame based legacy protocols (Ethernet and Token Ring), using only one ATM header
22 for up to 31 cells from the same virtual circuit in a packet. The specification of CIF over PPP
and Sonet is underway. A significant feature of CIF is that ATM can be transported to
24 workstations without changing the legacy NIC (network interface controller) card because the
necessary processing is done in simple downloaded software "SHIM" on the workstation.

2 LAN Emulation (LANE)

4 LAN Emulation is a technique that specifies the interfaces and protocols
needed for providing LAN-supported functionality and connectivity in an ATM environment,
6 so that legacy protocols can be interoperable with the ATM protocols, interfaces, and devices.

8 In legacy LANS, the membership of an individual station to a LAN segment is
dictated by the physical connection of the station to the physical shared medium.
10 Membership of a station to an ATM LAN segment is identified by logical connections to the
multicast ATM virtual connection. Hence, membership of an ATM LAN segment is defined
12 logically rather than physically; the membership information is stored in some management
database. This capability of ATM LANs offers terminal portability and mobility. LANE
14 does not provide transparent support for LAN-based application since it functions at layer 2,
like a bridge. Effectively it is a converting-bridge technology between the connectionless
16 Ethernet/Token Ring environment and the connection-oriented ATM environment. It also
supports ATM-enabled devices to communicate with LAN Emulated devices.

18 LAN Emulation does not allow users to leverage the end-to-end class of
20 service functionality which ATM provides in end-systems; however, it will provide for a
higher bandwidth and a more stable network infrastructure for large building and campus
22 backbones.

Multiprotocol Over ATM (MPOA)

2

MPOA is a set of standards designed to support distributed routing protocols
4 other than IP. The functionality is developed on top of LANE and NHRP (Next HOP
Resolution Protocol, a protocol proposed for ATM address resolution based on classical IP).
6 MPOA can be viewed as solving the problems of establishing connections between pairs of
hosts that cross administrative domains, and enabling applications to make use of a network's
8 ability to provide guaranteed quality of service. Provided below is a summary comparison of
LANE and MPOA, including advantages of MPOA and benefits of MPOA:

10

MPOA versus LANE

12

MPOA is an evolution of the LAN Emulation model. MPOA will make use of
14 the LAN Emulation services.

16

LAN Emulation operates at OSI layer 2, hence, it's bridging.

18

MPOA operates at both OSI layer 2 and layer 3, hence, it is both bridging and
routing.

20

LAN Emulation hides ATM/QoS, MPOA exposes both.

22

Advantages of MPOA

2

Clients can establish direct connections to remote servers without using

4 routers

6

Lower latency in establishing connections between devices

8

Reduced amount of broadcast traffic

10

Flexibility in selection of Maximum Transfer Unit size to optimize

performance

12

Benefits of MPOA

14

It provides the connectivity of a fully routed environment

16

It takes advantage of ATM, direct interdomain connection, and QoS

18

It separates switching from routing

20

It provides a unified approach to layer 3 protocols over ATM

22

Multiprotocol Transport Networking (MPTN)

24

MPTN has its roots in a multivendor, multiprotocol networking model developed by IBM in 1992 called the Networking Blueprint. Presenting a somewhat different view than that of the OSI reference model which describes a single way to implement networking technologies, the Networking Blueprint described a way for a number of unlike networking technologies to coexist. In 1994 the Networking Blueprint was expanded (also renamed the Open Blueprint), by opening up the application section turning it into a model for networking as well as a structure for a distributed systems environment in which distributed applications can run.

The Open Blueprint can be separated into four areas (Applications, Application Enabling Support, Distributed System Services, and Network Services). Although this research can draw from all sections of the model, there is particular interest in Common Transport Semantics (CTS) which separates Distributed Services from the transport network layer of Network Services. CTS is an important section in the Open Blueprint because it provides the place where the multiprotocol transport architecture can be implemented. It provides a place where a set of transport semantics common to all transport network protocols are provided. This means that the applications in the top section, using their respective APIs and communication programming styles, can select and work with any transport network, regardless of the communications protocol the transport network implements. CTS, therefore, provides a way to separate the APIs from their original transport networks, allowing them to run on top of other types of transport networks. When the protocols do not match, CTS becomes the glue between them. CTS bridges the gap between the needs of the user of the transport network and the services provided by the underlying transport network itself.

2

Virtual LAN

4

Two slightly different views have been presented of a virtual LAN. One view focuses on the virtual LAN as a networking environment where users on physically independent LANs are interconnected in such a way that it appears as if they are on the same LAN workgroup. A second view focuses on the concept of a virtual LAN being a partitioning of one physical network into several logical networks. The reconciliation of these views comes from the observation that a network, as used by the first view may be composed of multiple physical LANs which is the beginning of the second view. Even though these are somewhat different views, aggregation and partitioning can be reconciled and are, in fact, talking about the same thing. There may, however, be greater utility in using one mode of viewing the problem over the other in certain situations.

14

The following observations regarding the two primary orientations (Port-centric and Device-centric) of virtual LAN models may be useful in later deliberations:

16

The port-centric model defines a virtual LAN as a collection of physical ports either associated with LAN or ATM switch interfaces. Clients are manually assigned to a virtual LAN, and the ports that make up the virtual LAN are kept in a database. This model operates as a MAC layer bridge transporting messages between members of the virtual LAN.

22

In the device-centric model, hosts are identified by either their MAC address or their Network Layer address. In the device-centric model network, the administrator assigns clients to a virtual LAN group using their address as an identifier. Both the location

24

of clients and the transport of data between clients are managed by the network and are based
2 on the client's address.

4 **Virtual Network**

6 Virtual networks are somewhat less clearly defined than virtual LANS. The
use of the term virtual network seems to be more closely aligned with various telephone
8 service providers. In this context, the term is used to describe a logically defined network
overlaying the physical network belonging to the carrier. The notions in this description are
10 analogous to the second view relating to virtual LANS. In this research, however, notions
more in line with the first view of virtual LANs will be used. From this perspective, our
12 major interest is in dynamically connecting and disconnecting network domains that may or
may not belong to the same organization. The process of dynamically connecting and
14 disconnecting then merging and separating logical networks is the process referred to in this
work as creating/destroying virtual networks, and the connected state is referred to as a virtual
16 network.

18 **QoS Guarantees**

20 The significance of time with regard to the present invention is that it
correlates with the need for quality of service, or QoS, guarantees from the network
22 connecting the collaborators. The closer to real-time interaction that the collaborative activity
requires, the greater the need for QoS guarantees from the network. Even collaborative
24 content such as voice or video, when they can be delivered at a later time (e.g., voice mail
messages and archived video segments), does not require QoS guarantees from the network.

The fact that they can be shipped to the corresponding collaborator for use at a later time essentially removes the need for QoS. Sometimes the nature of the collaborative content (e.g., large graphics files) is more effectively handled by higher bandwidth network links, but this is still not the same as requiring QoS guarantees from the network.

Service Template

A service template is used herein to describe the set of characteristics that is needed to set up a particular telecommunications service. A service template may be predefined and may be accessed by agents within the participating network to determine the attributes and parameters needed to set up and execute a particular telecommunications service. For example, the user may use a whiteboard template to set up a whiteboard.

Control Architecture and Mechanisms

Having described some definitions used herein, a description of control architectures and mechanisms is presented below. With regard to control architectures and mechanisms, the taxonomy of distributed systems serves as a useful starting point. The following four dimensions of classification for categorizing systems may be used:

- Hierarchical and Peer-to-Peer
- Hot and Cool
- Tight and Loose

2

• Non-Redundant versus Replicated

4

Drawing on these dimensions, a taxonomy may be developed which provides some insights into the relatedness of various kinds of distributed systems and into the formation of the four primary categories used for discussion purposes (i.e., Message Passing, Client/Server, Distributed Database, and Distributed Transaction Processing).

8

Literature Considered

10

Prior to the present invention others have attempted to develop systems that would extend core networking more directly into the world of applications. In order to help organize and filter the literature that has been reviewed, a simple system of 10 categories within four general areas of interest was created. The general areas (Environment, Services, Networks and Implementation) seemed to provide an adequate characterization of the very broad range of topics that are intertwined with the question of orthogonal control of network resources.

18

Overview

20

The actual categories used for review are grouped within these four broad areas in the following manner:

24

Environment

- 2 1) Distributed Computing
- 4 2) Real-time Systems
- 6 3) Dynamic Programming
- 8 Services
- 10 4) Service Creation
- 12 5) Quality of Service
- 14 Networks
- 16 6) Network Control
- 18 7) Network Architecture
- 20 Implementation
- 22 8) Agent Based Software
- 24 9) Hardware Accelerators

2 10) Protocols

4 The research interests characterized by these categories are summarized in the
following sections which attempt to clarify the intent of the category as well as indicate some
6 of the topical issues identified by the literature review process.

8 **Distributed Computing**

10 Implicit in the research problem is a series of questions and issues as to what
constitutes distributed computing in a collaborative peer-oriented environment as well as
12 what are the best ways to implement such facilities. Systems with such characteristics are
actually a specialized subset of distributed computing systems. In this inquiry, there is
14 interest in both application frameworks and the supporting infrastructure constructs for such
distributed environments.

16 Arnold, Bond and Chilvers (1997) provided a useful introduction to the topic
of distributed-processing environments with the presentation of their framework Hector.
18 Starting with a discussion of the international standard for such environments known as the
"Reference Model for Open Distributed Processing" (RM-ODP), they compared several
20 prominent frameworks (APM's ANSAware, OSF's DCE, the OMG's CORBA, and
Microsoft's DCOM) with RM-ODP. Their examination pointed out that not only do these
22 prominent frameworks offer varying levels of the functionality specified by RM-ODP but
they are also unable to interoperate among themselves. These researchers stated that the work
24 on Hector was designed to provide "a framework that sits above other distributed

environments, providing open negotiation and interoperability of communication protocols,
2 high-level description of component services and their requirements, a rich set of support
services for objects, and an interaction framework that allows the description of workflow-
4 like interactions between autonomous objects." The very nature of the objectives of this
research provided valuable insights into the problems and issues of distributed-processing
6 environments. From a more pragmatic perspective, however, the majority of literature
covering research specifically targeted at aspects of distributed network infrastructure used
8 the CORBA framework [Vinoski (1997), Schmidt, Gokhale, Harrison, and Parulkar (1997),
Crane and Dulay (1997)].

10

There are several supporting constructs or mechanisms that are significant in
12 considering the implementation of infrastructures for distributed computing. Genereaux
(1997) described the InterLanguage Unification System (ILU) as a system "designed at Xerox
14 PARC with the purpose of providing a coherent model for distributing applications and
components across machine and network boundaries." According to Genereaux, ILU was
16 mainly concerned with defining interfaces between collections of program units called
modules which "can be written in different application languages, can exist on separate
18 machines, or can be distributed systems implemented by many program instances on many
machines. " Other useful mechanisms included the configurable event service for distributed
20 systems described by Mansouri-Samani and Sloman (1996) and the private access channel
security mechanism for shared distributed objects described by Dollimore and Xu (1993).
22 Shatz (1984) provided a survey of communication mechanisms for programming distributed
systems which was also useful as an introduction to the basic mechanisms and techniques
24 supporting distributed computing.

2 Beyond the questions involved with distributed application frameworks and
specific mechanisms, there remains the foundation issues associated with the basic paradigm
4 used in developing the distributed environment. Because of the distributed nature of the
communication infrastructure environment, message passing will necessarily become part of
6 the solution at some level of implementation. Initially, however, the message-passing model
is set aside in favor of the Linda type model for providing the framework for the solution
8 being sought through this research. The Linda model may be described as a coordination
mechanism which is orthogonal to a base language according to Carriero and Gelernter
10 (1989). Even though Linda is a very distant second to message-passing models in terms of
total research effort, as gauged by the frequency of its occurrence in the literature, it has a
12 number of properties that recommend it for the problem at hand.

14 The desired attributes of loose coupling and decentralized peer-oriented
control can be expressed quite easily in Linda-type constructs. According to Gelernter
16 (1985), Linda is fully distributed in space and fully distributed in time. Linda programs are
collections of ordered tuples which may contain either executable code or passive data values.
18 "The abstract computation environment called 'tuple space' is the basis of Linda's model of
communication. An executing Linda program is regarded as occupying an environment
20 called tuple space or TS. However many concurrent processes make up a distributed program,
all are encompassed within one TS."

22

 The communication used by Linda-type programs is based on generative
24 communication which has the characteristic of communication orthogonality in which the
creating and consuming processes have no knowledge of one another. Gelernter stated that

"communication orthogonality has two important consequences: space-uncoupling (also referred to as 'distributed naming') and time-uncoupling. A third property, distributed sharing, is a consequence of the first two." Space-uncoupling referred to the fact that a tuple in TS may be input by any number of disjoint address-space processes. Time-uncoupling referred to the situation that a tuple added to TS remains until it is specifically removed. Finally, distributed sharing allowed disjoint address-space processes to share some variable by depositing it in TS.

8

The notion of a single logical tuple space creates implementation issues when the logical single tuple space is in actuality composed of multiple disjoint tuple spaces that must be explicitly synchronized to form the logically unified tuple space. Some of the issues involved are addressed in research into this topic by Kambhatla (1991) in his Masters Thesis dealing with replication issues for distributed and highly available Linda spaces. Though not developed in conjunction with the Linda model, Thekkath (1994) in his Ph.D. dissertation, regarding system support for efficient network communication, supplied considerable insight into new models of fast efficient communication that are particularly well suited to the ATM environment and provide basic mechanisms for the maintenance of shared global memory.

18

Returning to the suitability of the Linda model (assuming the presence of reliable distributed shared global memory), there were several Ph.D. dissertations and research reports which were available and could form the formal underpinnings of a solution based on the Linda model. This research, however, seeks to expand the possible range of solutions by considering the possibility of allowing the multiple tuple spaces that make up the single logical tuple space to only be loosely synchronized and to rely on smart agents at each of the separate nodes to detect and respond to perceived inconsistencies. This relaxed view of

the nature of tuple space should simplify many of the complexities that accompany use of the Linda programming model in distributed applications. It does, however, mean that the programming model will only be Linda-like and not the Linda programming model as defined. There was a collection of useful papers dealing with Linda-like systems and their implementation that was assembled at the Edinburg Parallel Computing Centre. These papers described systems and implementations that alter the basic Linda model for various reasons. These papers contributed a number of useful insights into the possible use of only parts of the original Linda model [Bal, Kaashoek, and Tanenbaum (1991); Broadbery and Playford (1991); Butcher (1991); Callsen, Cheng, and Hagen (1991); Carriero and Gelernter (1991); Schoinas (1991), Schouten and van Nieuwkerk (1991); Thomas (1991); Wilson (1991); Zenith (1991)].

Real-time Systems

From the perspective of collaborative applications, Real-time is a somewhat subjective notion that refers to the ability of the supporting infrastructure to respond to the supported activities and applications at a pace that does not retard the natural time progression normally associated with the activities being supported. An introduction to the basic concepts and issues of real-time communication was provided by Verissimo (1993). Beyond the basic concepts and techniques of analysis, design, and implementation, this research had some specific interest in various mechanisms that would be suitable for use in the target network environment. Works such as that by Schmidt, Gokhale, Harrison, and Parulkar (1997) on an end system architecture for real-time CORBA were certainly of interest; however, there were even more fundamental questions that were raised by Kopets and Grunsteidt (1994) in their work on TTP, a protocol for fault-tolerant real-time systems.

2 Kopets outlined the two fundamentally different paradigms for the design of
real-time systems (i.e., event-triggered architectures and time-triggered architectures). After
4 discussing the advantages and disadvantages of each, he proceeded to describe his Time-
Triggered Protocol (TTP). Even though this protocol was initially designed for automotive
6 control applications, there were a number of features that were worthy of consideration for
the coordinating control of distributed agents that might be used to implement the distributed
8 service creation facility that is an objective of this research. This research area was of
considerable importance in finding an implementation strategy that is operationally viable
10 and scalable across private and public networks that range from small to global in size.

12 **Dynamic Programming**

14 The current interest in dynamic programming arises primarily from issues
related to naming, data structures and binding. Helary and Raynal (1988) developed
16 algorithms for assigning distinct identities to sites of an anonymous distributed system. This
was an issue of considerable concern when constructing a dynamic distributed control and
18 service creation facility that needs to scale to large networks and to encompass multiple and
varying collections of domains associated with both public and private networks.

20 The proposed environment of the present invention is itself an application or
22 set of applications. From this perspective, there was interest in the work of Warren and
Sommerville (1996) which presented a model for dynamic configuration while preserving
24 application integrity. There is, however, a heightened interest in works which address issues

associated with dynamic linking and modification of software through on-the-fly software

2 replacement.

4 The research basis for this interest stems from implications of the intended
distributed control and service creation architecture. In the centralized or client/server model
6 for control and service creation, there is a well defined and more manageable group of system
elements that allow for scheduled downtime for maintenance. In a distributed peer-oriented
8 model, the situation is somewhat different. Since control and service creation functionality
would run on workstations as well as network devices, the systematic system wide
10 maintenance and enhancement of control and service creation functionality is no longer a
viable option. Furthermore, because of the dynamic membership of service creating groups,
12 there is an increased risk of non-interoperability due to conflicts among versions. This
problem would need to be addressed in a manner that could be used in large networks that
14 possibly are owned by different entities.

16 One approach to solving this architectural problem could use the capability of
automatically synchronizing version levels whenever necessary. Such a capability would
18 most likely need the ability to perform some sort of on-line maintenance. One of the most
relevant pieces of research came from Hauptmann and Wasel (1996) in which the researchers
20 described specific techniques for accomplishing on-line maintenance by using on-the-fly
replacement of software components. Cowan, Autrey, Krasic, Pu, and Walpole (1996) as
22 well as Veitch and Hutchinson (1996) contributed valuable concepts with their works dealing
with adaptive, extensible and configurable operating systems. Queichek's (1996) description
24 of dynamic configuration management also provided a source of issues and ideas relevant to
this problem area. Finally, Virtual Computer Corporation (nd) provided a useful overview of

reconfigurable computing. Although their perspective was hardware oriented, some of the

2 concepts appeared to be useful.

Service Creation

4

The issue of service creation is a central concern to this research which seeks
6 methods and mechanisms to perform this function in a distributed and collaborative manner.

In reviewing the literature associated with this research area, the following three categories
8 will be used to help focus on aspects that are particularly important to this work:

10

Object-Oriented Approach to Services

12

Service Models and Mechanisms

14

Multi-party Interactive Multimedia

16

A distributed object-oriented approach to defining and implementing
network based collaborative services holds great appeal from the standpoint of managing the
18 inherent complexity associated with such functionality. A direct introduction to this approach
was provided by Amouat, Brossard, Louvet, and Risser (1991) in their research into the
20 application of the object-oriented paradigm to modeling telecommunication services. Further
insight into specific techniques for dealing with services and service creation were provided
22 by Takura and Ohta (1994) in their research dealing with the generation of
telecommunication software from service specifications in state transition models. While this
24 work was not explicitly object-oriented, the content can certainly be applied in such an
approach. Schmidt (1995), however, used an explicit object-oriented approach in his work on

design patterns for initializing network services. "The Acceptor and Connector patterns

2 described in this article decouple the passive and active establishment of a connection,
respectively, from the service performed once the connection is established." This work also
4 addressed forces due to additions by using asynchrony to actively establish connections with
a large number of peers efficiently.

6

There are many issues and topics that can be addressed under the heading of
8 service models and mechanisms, but few can be ascribed the practical level of importance as
the topic discussed by Crowcroft, Wang, Smith, and Adams (1995). This article provided a
10 comparison of the Internet Engineering Task Force (IETF) and the ATM service models.
Some of the pragmatic interest comes from the fact that these service models are aligned with
12 the historically existing division of networks into two major categories: the Internet group
which has provided a best effort datagram service and the digital telephone networks which
14 have provided a reliable constant bit rate circuit service. Advances in networking and
computing have advanced packet/cell switching to a point where it is appropriate for use in
16 delivering a range of services. It is a logical next step, based upon economics, that these
services should be provided at a single access point to all end systems. This evolution has set
18 up an escalating confrontation between the Internet approach, proposed by the data
networking group, and the ATM approach, proposed by the telephony networking group.
20 ATM with its origins in telephony has a massive investment in developing standards for
reliability, accountability and quality of service (QoS). The Internet group has been working
22 to extend its model to accommodate more reliability and accountability. QoS has been
approached by the Internet group by adding an optional resource reservation protocol called
24 RSVP.

2 The issue over which approach will gain dominance is very significant to the
research that is proposed. The ATM approach is inherent in the design of ATM, occurs at
4 lower levels of the protocol stack, and is established during call setup. The RSVP approach is
optional. It allows applications that have QoS needs to communicate them to the network
6 while allowing applications that don't to function the same way that they always have.
According to Crowcroft et al. (1995), "RSVP is not a call setup; it is based on the receiver's
8 signaling of requirements, rather than an initiator for two reasons:

10 Multicast, or many-to-many communicating applications, can rendezvous with
a signaling complexity order (constant), rather than order (n^2).

12 Receivers can signal different reception qualities, independent of the senders'
14 selection of quality."

16 Lambert (1997), writing for an industry trade magazine, reported on the recent
attempts at reconciliation between the two opposing approaches. According to Lambert,
18 there were new proposals introduced in mid-1997 which addressed the QoS issue at the
transport, network and media access layers as well as the application layer (layer 4). Lambert
20 reported that ATM manufacturers considered mapping RSVP to the QoS mechanisms of
ATM as the next natural step following the merging of the router and switch. From the IETF
22 side, there is a working group called the Integrated Services Over Specific Lower Layers
(ISSL) Group that is also taking up the challenge of integrating the Layer-4 RSVP
24 mechanisms used within Ethernet networks to ATM QoS mechanisms used at Layers 2 and 3.

2 As significant as this issue of reconciliation between RSVP and ATM is, there
are other topics that are of interest in the context of the proposed research. Bocking (1996)
4 described a uniform application programming interface for basic-level communication
services. Poo and Chew (1996) described the modeling of the XOM/XMP application
6 programming interface (API) which gives access to the services of common management
information service (CMIS) and the Simple Network Management Protocol (SNMP) XOM
8 API provides a general-purpose data handling mechanism and XMP API provides service
primitives to network management protocols, both of which are useful in the service creation
10 effort. Other mechanism-related works that are of interest included Parris, Ventre, and Zhang
(1993) who dealt with the graceful adaptation of guaranteed performance service connections
12 and Ferrari, Gupta, and Ventre (1995) who discussed issues related to distributed advance
reservation of real-time connections.

14

 The final review category, multi-party interactive multimedia (MIM)
16 addresses issues involved with supporting such capabilities which are the focus of research
interest. Szyperski and Ventre (1993) proposed a solution to efficient group communication
18 with guaranteed QoS based on an abstraction called the Half-Duplex Real-Time Channel.
This abstraction, the researchers asserted, "reduces the complexity of the creation of network
20 support for common MIM applications, and decreases the amount of resources to be reserved
in the network." Effelsberg and Muller-Menrad (1993) and Moran and Gusella (1992)
22 contributed ideas regarding dynamic multiparty support for applications. Gupta et al (1994)
took a slightly different perspective in discussing a scalable resource reservation for multi-
24 party real-time communication.

2

Quality of Service

4

A great many of the challenges and issues involved in service creation relate to the quality of service (QoS) and the methods of requesting and obtaining specific QoS guarantees from the network. The literature reviewed in this section has particularly close ties to work reviewed in the section on service creation as well as on the upcoming section on network architecture. In fact, much of this literature could be included in the reviews of these other sections. They are, however, reviewed here because of the overriding focus on QoS issues. In reviewing the literature, three categories (architecture, mechanisms, and performance issues) seem useful in organizing the material.

12

In the area of architecture, Lakshman and Yavatkar (1996) described AQUA which is an adaptive end-system QoS architecture. AQUA provides a framework for managing resources such as CPU, network interface, memory, and bus bandwidth. The researchers asserted that the "significant and novel contributions of AQUA include an adaptation framework, QoS specification, resource managers, and an application level QoS manager that performs application-based graceful adaptation when resource requirements change or the demand for resources exceeds available capacity." A further contribution was a CPU management algorithm called RAP (Rate-based Adjustable Priority Scheduling) that provides predictable service and dynamic QoS control.

22

Other useful pieces of research include Moghe (1996) who described a new paradigm called "Enhanced Call" for applications with dynamic client-membership and client-level binding in ATM networks. The approach sought to guarantee an upper bound on

0397321.03

client-level degradation by statistically reserving virtual channel links for potential new client arrivals. Gopalakrishnan and Parulkar (1996) presented a framework for providing QoS guarantees within the end system for networked multimedia applications. The framework contained the following four components: QoS specification, QoS mapping, QoS enforcement, and protocol implementation. Also included in the framework was an application level protocol implementation model along with techniques to reduce the cost of data movement and context switching in such implementations.

Under the mechanism classification, the papers selected were either associated with protocols and client/network boundary issues or with aspects of flow control. Within the client/network category, there were three papers of particular interest. The first, by Campbell and Coulson (1997), described the implementation of an adaptive transport system that incorporates a QoS-oriented API and a range of QoS mechanisms that assist multimedia applications in adapting to fluctuations in the delivered network QoS. The second paper by Ferrari, Ramaekers and Ventre (1992) investigated the feasibility of providing an extended client interface that allows more flexibility in the client-network interactions. The researchers claimed that "the proposed model improves the utilization of network resources, and increases the network's capability to support multimedia traffic, while continuing to offer a guaranteed quality of service." The final paper by Lowery (1991) proposed a real-time delivery system composed of a new protocol for administration of real-time connections combined with modifications to the Internet Protocol (IP) to support such connections.

22

The flow control aspect was touched on by Ohnishi, Okada, and Kiyohiro (1988) who examined two mechanisms for providing performance and flow control management with respect to performance (QoS) control, introduction of delay, and loss

sensitive service classes. Zhang and Keshav (1991) examined six new queue service disciplines (Virtual Clock, Fair Queueing, Delay-Earliest-Due-Date, Jitter-Earliest-Due-Date, Stop-and-Go, and Hierarchical Round Robin) in order to show why each discipline can or cannot provide bandwidth, delay and delay jitter guarantees.

In the literature specifically dealing with performance, there was a quantitative study by Tobagi and Dalgic (1996) which focused on the performance of ethernet and ATM networks carrying multimedia traffic, particularly audio and video traffic. Also, Moldeklev and Gunningberg (1995) studied deadlock situations that can occur when using TCP over ATM. This last issue is of great practical importance given the necessity to interoperate with the great installed base of TCP and IP based networks.

Network Control

As stated earlier, one embodiment of the present invention is concerned with the structure of an application level or "logical level" control mechanism for communication services used in support of various peer-oriented types of applications. The focus of research for the present invention centered on the identification of a suitable infrastructure for component-based control of communication services which is orthogonal to the underlying native control mechanisms of the network being used. The application domain is one axis (the source) of this proposed orthogonal mapping relationship, while the literature reviewed in this section on network control and the following section on network architecture represent aspects of the target axis, i.e. the underlying network.

When dealing with multimedia traffic QoS issues at the network level, one of the critical topics of interest involves congestion control. Because this research is specifically interested in ATM networks, the papers by Lu(nd) and Jain (1995) provided a good foundation for the issues and approaches that are associated with this topic. The paper by Jain is particularly useful in that it provided a survey of the congestion control mechanisms for ATM networks that were selected by the ATM Forum traffic management group. Furthermore, Jain described the reasons why the adopted methods were selected from the available approaches and this insight was valuable in considering the current capabilities and direction for evolution of congestion control capabilities.

10

With the focus on dynamic service creation and ATM networks in particular, it is consistent that there would be considerable interest in the Available Bit Rate (ABR) service that has been defined for ATM. ABR is the newest ATM service category and according to Bonomi and Fendick (1995) it was designed to "systematically and dynamically allocate available bandwidth to users by controlling the flow of traffic with feedback." This service class has characteristics that would recommend it as a fundamental building block for the environment that this research seeks to establish. Chen, Liu, and Samalam (1996) described objectives of the new service as well as its relationship to other established ATM services and the existing agreements on the traffic control mechanisms to support it. Siu and Tzeng (1994) and Walthall (1995) provided other insights into the rate-based control framework that was adopted by the ATM Forum. Saito et al. (1996) examined the performance issues associated with a public ABR service, the support of TCP-over-ABR, and suggested a method for maintaining the throughput of point-to-multipoint ABR when the number of destination nodes is increased. These are all issues of great interest and the technical analysis approach used by these researchers was useful in providing a basis for reasoning about the use of this

service type in the proposed infrastructure. Fendick (1996) contributed with a discussion of
2 the evolution of controls for the ABR service with respect to the algorithms for determining
this feedback, an important topic since the details of how this is to be accomplished are
4 largely outside the scope of the ATM standards and specifications. Finally, Kolarov and
Ramamurthy (1995) discussed the performance implications of using the reactive rate
6 adaptation mechanism associated with ABR in wide-area networks. In this paper, the
researchers demonstrated that "the performance of virtual channels traversing large numbers
8 of hops in WANs can be substantially improved by giving priority to network transit traffic
over traffic entering the network."

10

Going beyond the broad topics of congestion control and the ABR service,
12 there were several papers dealing with specific mechanisms and techniques that were useful
as a starting point for considering mechanisms and approaches for the desired infrastructure.
14 Perhaps the most encompassing paper was provided by Lin (1997) which discussed the
operation, administration, and maintenance (OA&M) management for the Global System for
16 Mobile Communications (GSM). Although GSM is a European wireless digital signaling
network standard, it provides a common set of compatible services and capabilities that are
18 worth considering when framing a context for establishing a collaborative communication
environment such as the one in an embodiment of the present invention. The paper described
20 how the telecommunication management network (TMN) concept is applied to the GSM
OA&M.

22

Performance management, network traffic, and congestion control were topics
24 investigated in the following three papers. Gaiti and Pujolle (1996) sought "to introduce
performance monitoring aspects of asynchronous transfer mode (ATM) networks and then to